Return To Agenda Page

SPARC64™ VI / VI + Next Generation Processor

October 25, 2005

Takumi Maruyama

Enterprise Server Development Division,

Fujitsu Limited

SPARC64™ Processor Roadmap

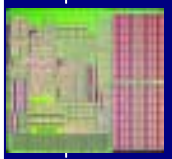

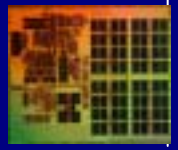

SPARC64™ V
1.35GHz
130nm

SPARC64™ V +
1.89GHz
90nm

2.16GHz
90nm

SPARC64™ VI
2.4GHz
90nm
2-core x 2-strand

SPARC64™ VI +
>2.7GHz
65nm
4-core x 2-strand

Multi-Threaded Processor

Mainframe RAS

2003 2004 2005 2006 2007 2008

– This road map is a previous notice and might change.

SPARC64™ VI / VI +

SPARC64™ VI Chip

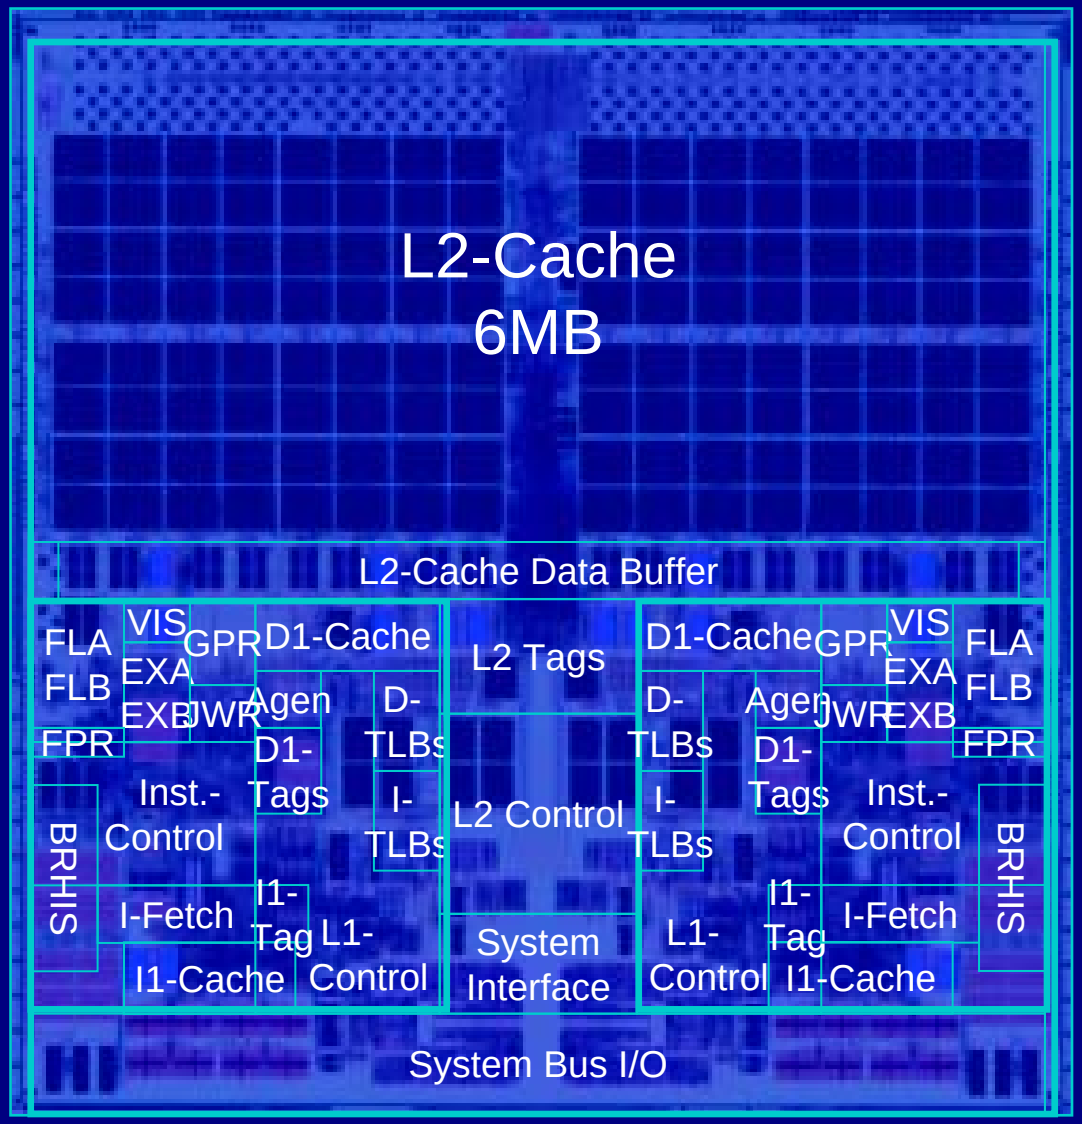

- Architecture Features
  - Multi-Threaded Processor
    2core x 2strands
  - Large L2$ and TLB
  - 2.4GHz
  - Jupiter Bus
- 90nm Cu 10 layers
  - 20.38mm x 20.67mm
  - 540M transistors
  - 412 signal pins
- 120 W

SPARC64™ VI MT Architecture

- Design Philosophy
  - Hide memory latency crucial to performance
  - Reuse proven SPARC64™ V core
    i.e., No major pipeline changes
  - Evolution rather than revolution to keep step with the SW environment
  - No compromise on single thread performance

➔ 2 core x 2 strand per Chip

➔ Switch strands on events: Vertical Multi Threading

Vertical Multi-Threading

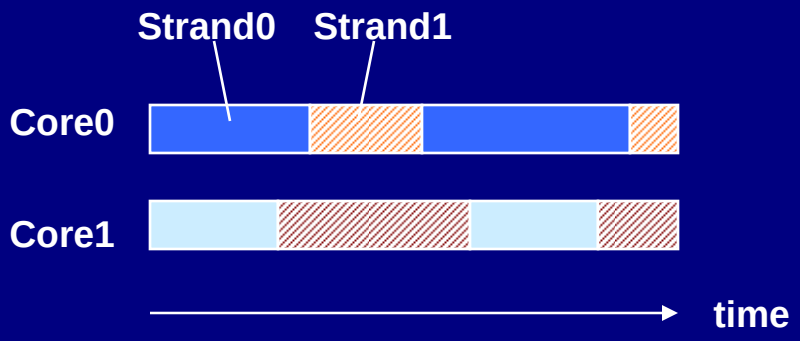

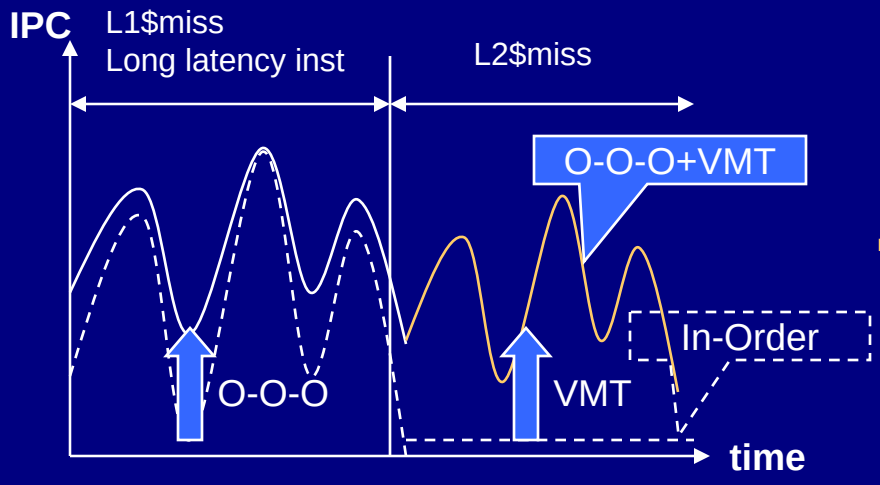

- Coarse Grain Multi-thread
- Strand Switch on events
  - L2$miss
  - Hardware Timer
  - Interrupt
  - MT control instructions

➔How to hide latencies

- L1$miss, Long Latency inst: O-O-O
- L2$miss: VMT

Pipeline structure of SPARC64™ VI

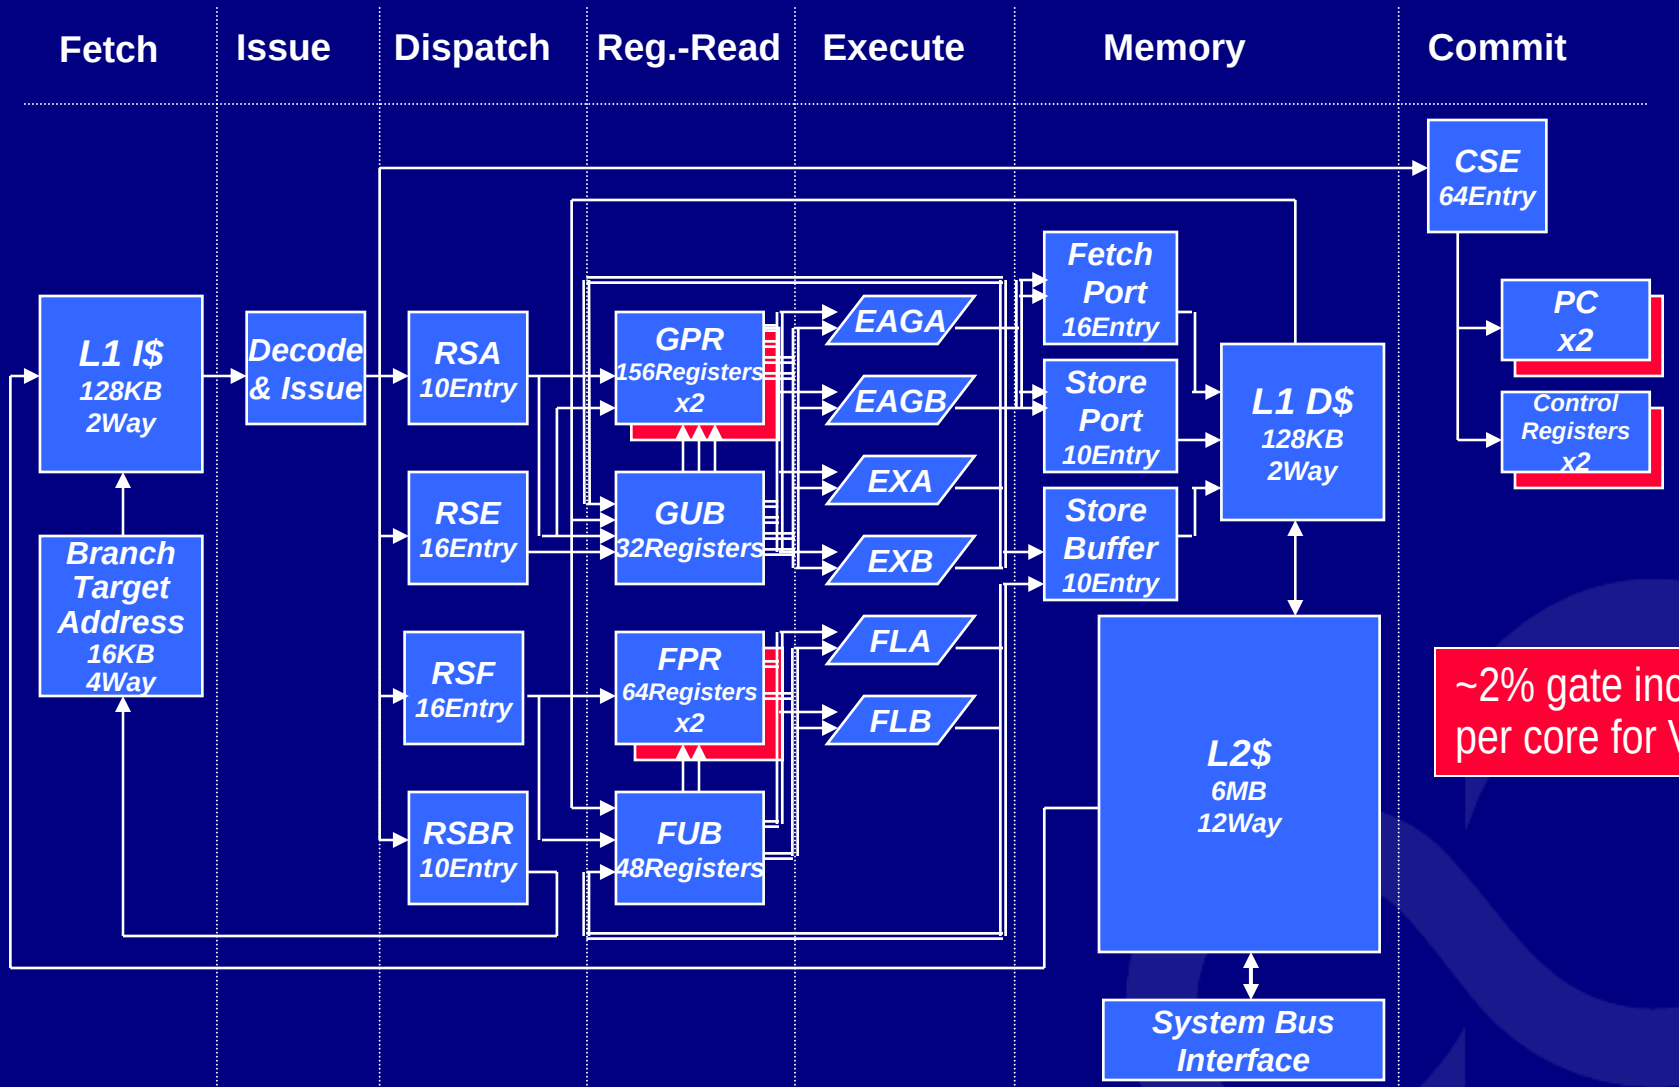

~2% gate increase per core for VMT

SPARC64™ VI / VI +

VMT Implementation Issues

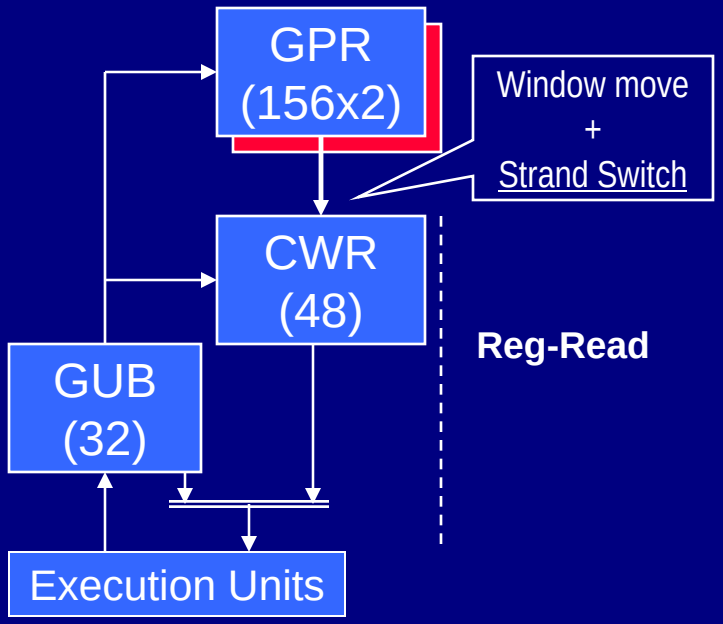

- SPARC-V9 Register Window
  - GPR
    - SPARC-V9 Integer Reg
    - 156reg(8win) x 2strand, 1R2W
  - Current Window Register
    - Subset of GPR, 48reg(2win), 8R2W

➔Keep the Reg-Read in 1-cycle

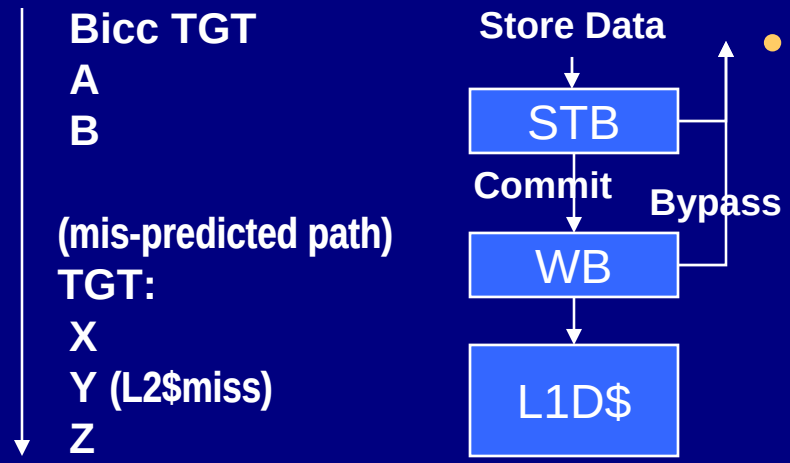

- Strand switch due to L2$miss
  - Speculative v.s. Non-speculative
  - Load v.s. Store*

    * Store L2$miss is hidden by SP & WB

➔Non-speculative, Load only

VMT performance

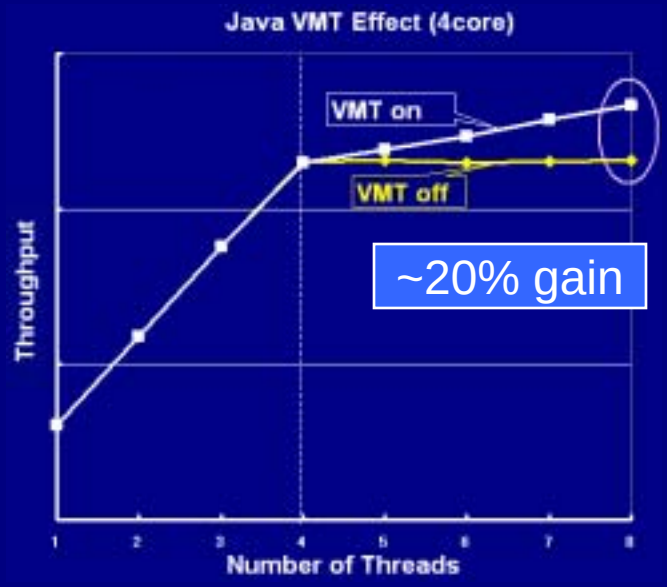

- Idle strands are "suspended" by the MT control instructions
  - VMT on = VMT off when Num of threads < cores
  - One strand can use all the resources (rename reg, reservation station...)
- VMT gives <u>additional</u> performance when Num of threads > cores

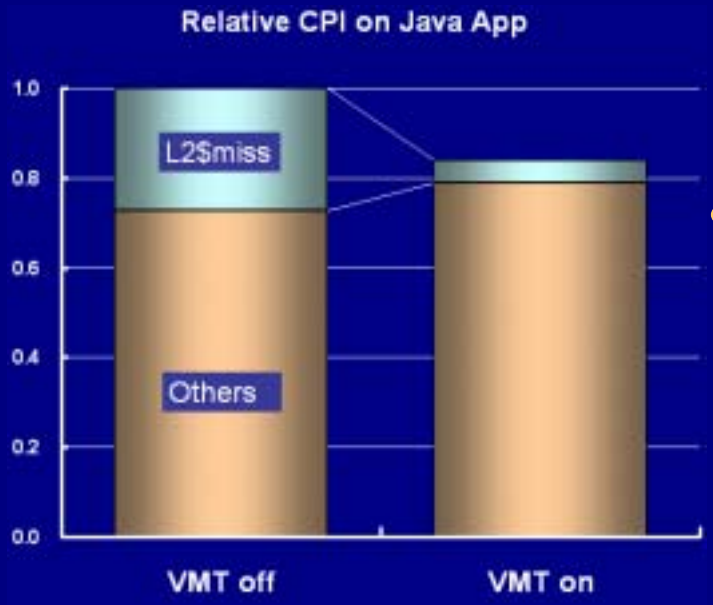

- L2$miss is hidden effectively with VMT

SPARC64™ VI / VI +

Other performance improvements

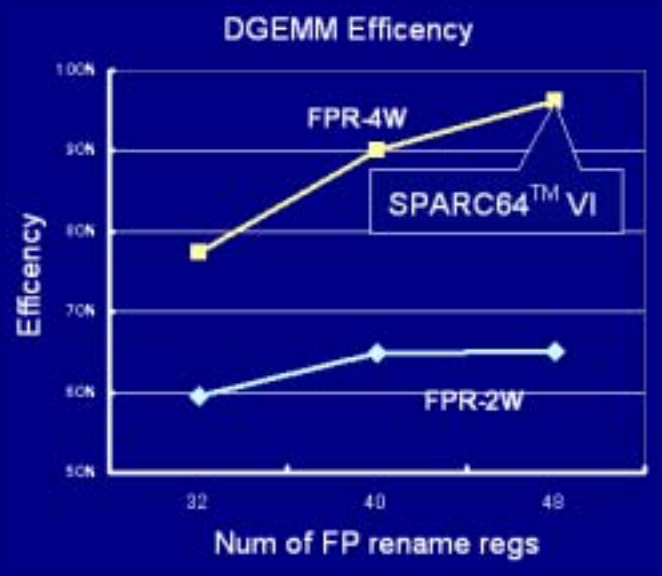

- Refined CORE
  - Fast FMA: 7 cycles
  - Rich FP rename Reg: 48
  - FPR: 4 Write-Port

  ➔Effective on DGEMM (matrix multiply)

  - Write-buffer
  - Doubled TLB
  - etc...

➔About **2x Single thread performance** of SPARC64V/1.35GHz

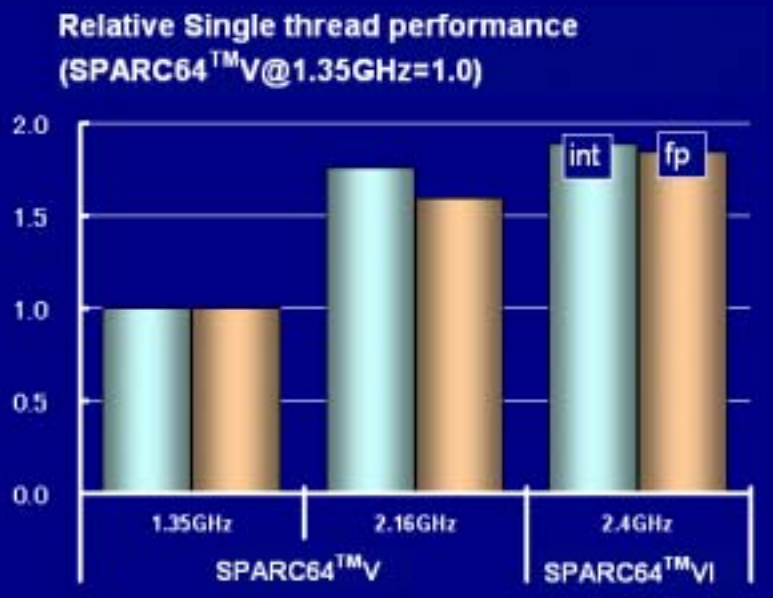

SPARC64™ VI / VI +

Reliability, Availability, Serviceability

		SPARC64™ VI
Cache protection	Tag	ECC (L1$: Duplicated & Parity)
	Data	ECC
Cache dynamic degradation		Yes
ALU / Register		Parity
HW Instruction Retry		Yes
History		Yes

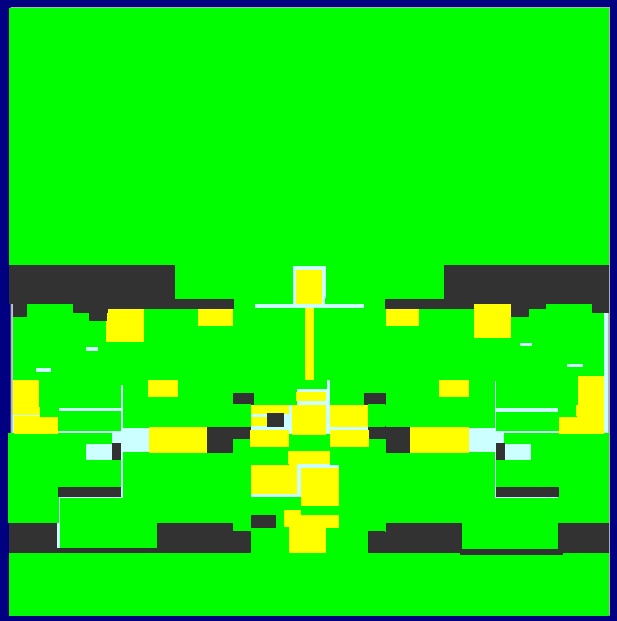

- New RAS features of SPARC64™VI
  - Granularity of CPU degradation: from chip to core
  - Number of checkers
    SPARC64™V: ~800
    SPARC64™VI: ~2,200

- RAS Coverage
  - Green: 1bit error Correctable
  - Yellow: 1bit error Detectable
  - Gray: 1bit error harmless
- All RAMs are ECC protected or Duplicated
- Most latches are parity protected

➔Guaranteed Data Integrity

Performance Evaluation

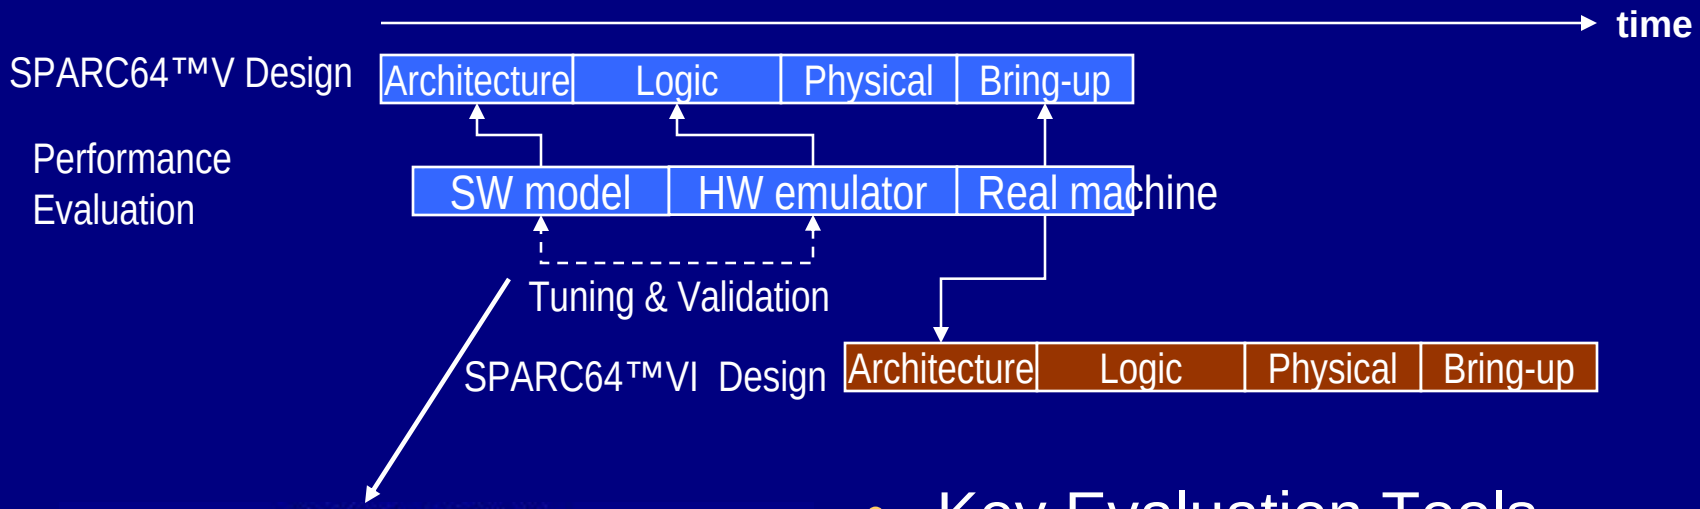

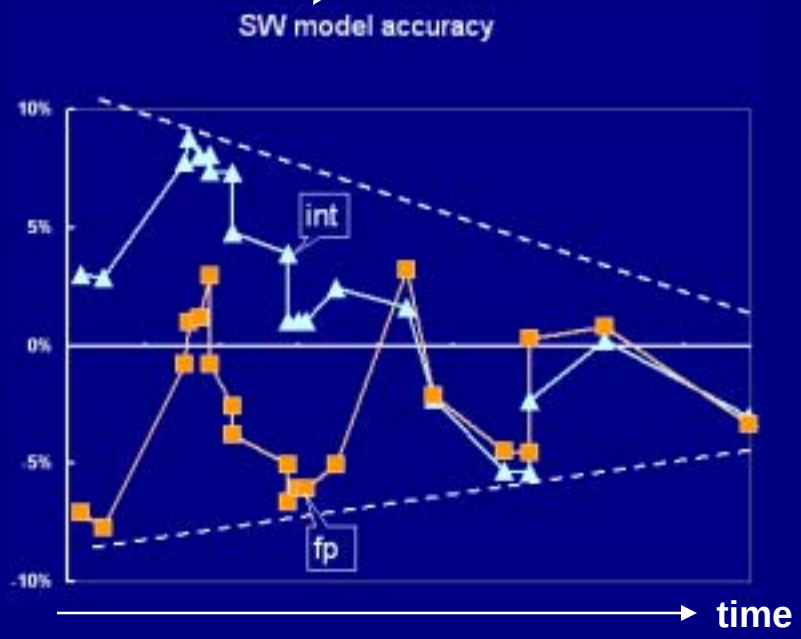

- Key Evaluation Tools
  - SW Performance Model
  - HW Emulator
  - Real Machine with OPSR
- Performance Accuracy
  - Improve as SW model matures
  - The final number of SPARC64™V: 4.2%(int) / 3.9%(fp)

SW model and HW Emulator

Performance Evaluation Flow

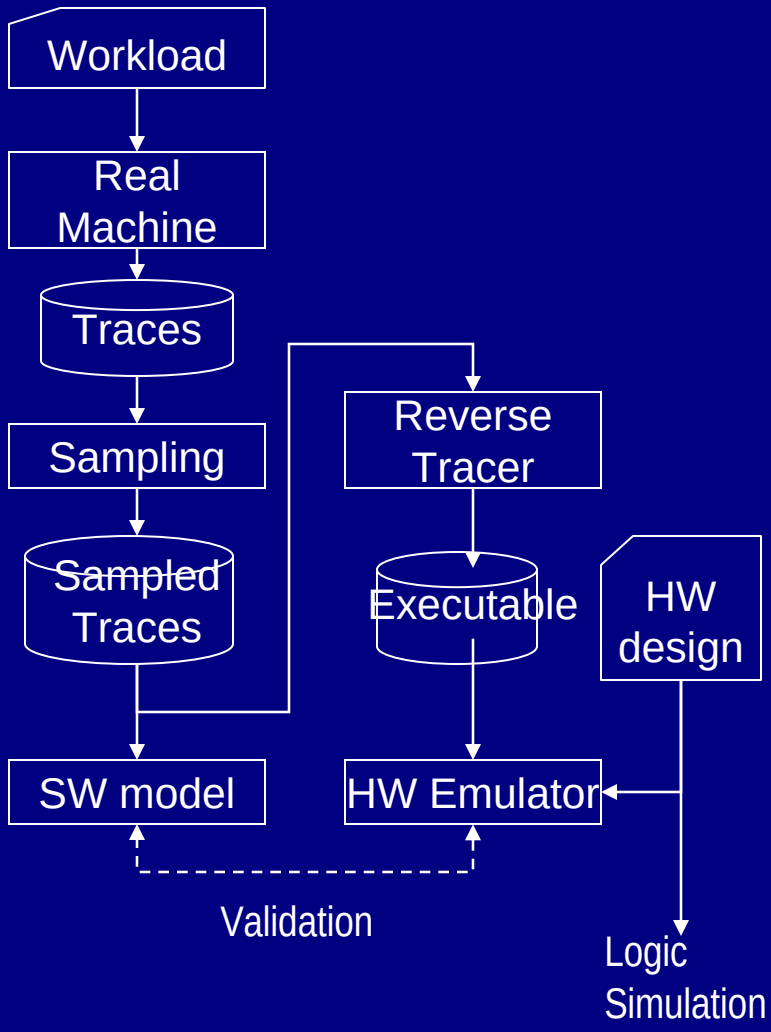

- Sampling
  - Find the best representative
  - Characteristic: CPI, $miss-rate, ...
- Reverse Tracer
  - Originally developed for Mainframe
  - Generate executable from a trace
  - Applicable to general Workload
- Execution Time (164.gzip)

	Number of Instructions	Execution Time
Real Machine (SPARC64™V/2.16GHz)	310G	122s
HW Emulator	1G	2.5h
SW Model	540M	10h

SPARC64™ VI / VI +

OPSR

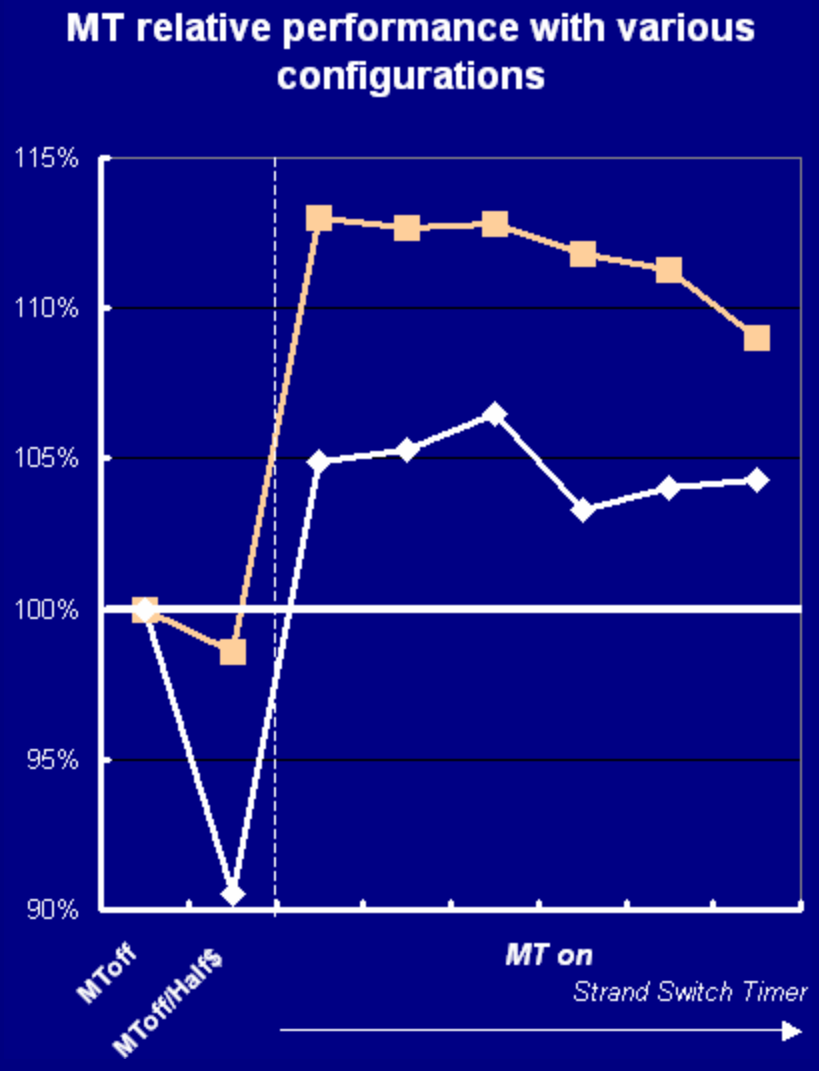

- Operating Status Register: Mode register to change CPU's configuration
- Objective
  - Find the best parameters
  - Gather information for future generation
- OPSR of SPARC64™ VI
  - Total: 1450 bits
  - VMT related: 170 bits

SPARC64™ VI / VI +

SPARC64™ VI + Chip

SPARC64™ VI + Block Diagram

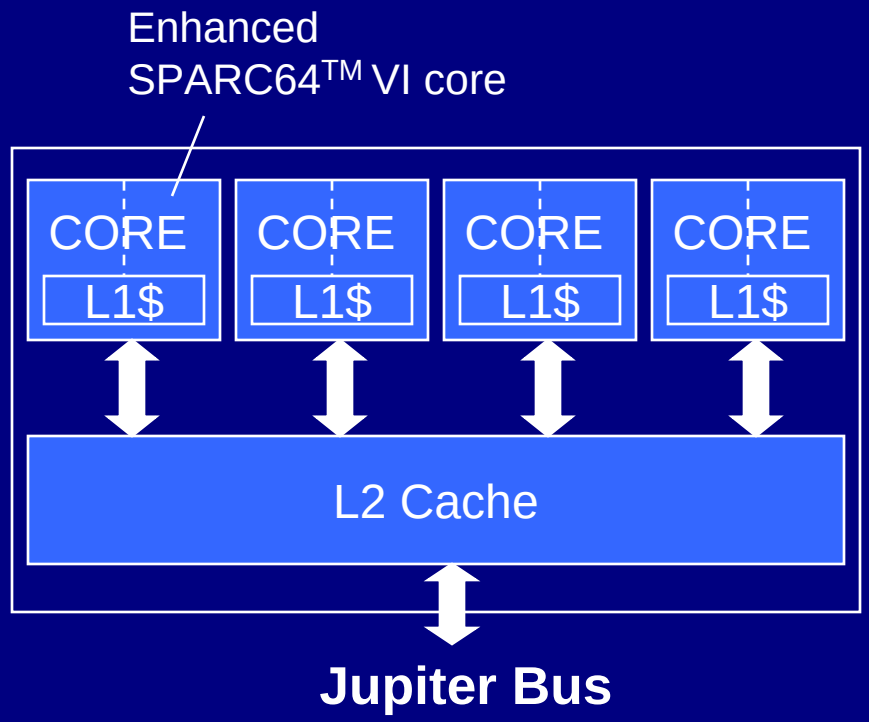

- Enhancements
  - Multithreaded processor 4core x 2strands
  - >2.7GHz
- Same System I/F: Jupiter Bus
- 65nm CMOS
  - 21.8mm x 21.3mm (TBD)

SPARC64™ Future Directions

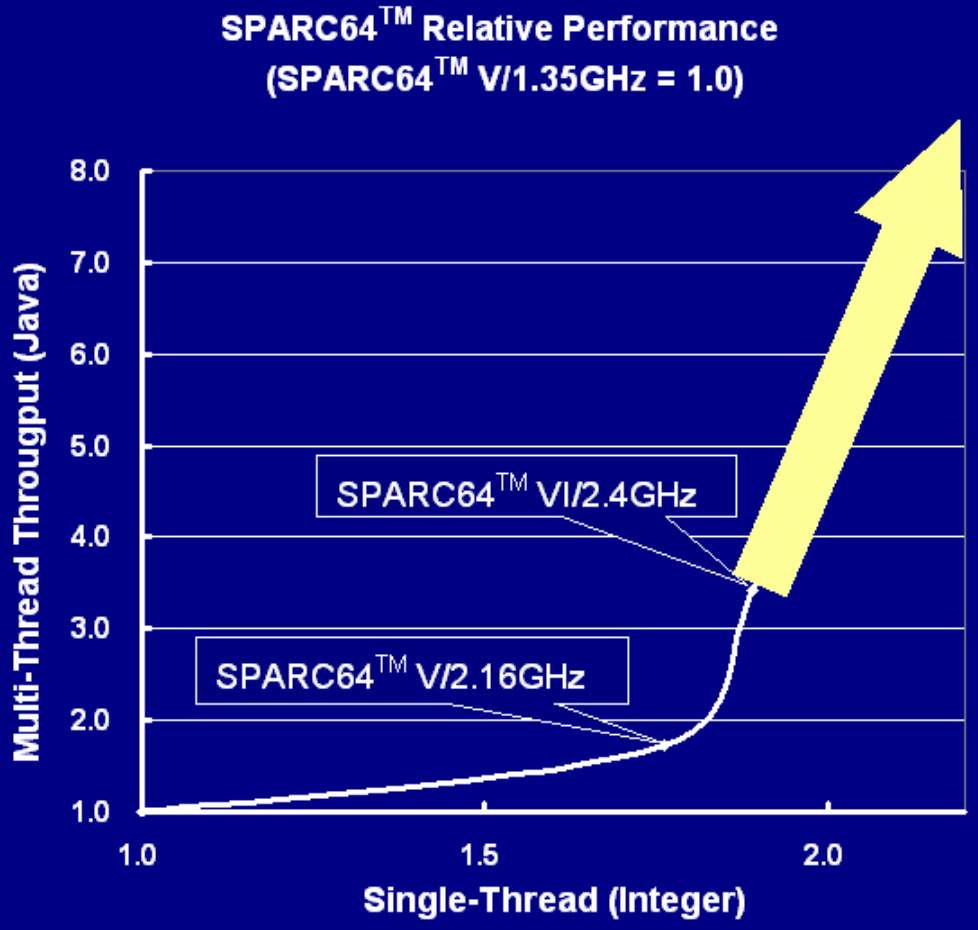

- Keep strong single thread performance
- Increase throughput performance with MTP (Multi-Threaded Processor)
  - multiple core
  - multiple strands
- Continuously develop to meet the needs of a new era

SPARC64™ VI / VI +

Summary

- SPARC64™ VI = SPARC64™ V plus
  - Higher Single Thread Performance
  - Higher Throughput: 2core x 2VMT
- O-O-O + VMT realizes high throughput with minimum design impact
- Evolutional development is a key to success of SPARC64™ series.
- SPARC64™ performance keeps growing as CPU of Fujitsu's PRIMEPOWER™ UNIX® server